

TECHNICAL SKILLS	<ul style="list-style-type: none">• Languages: Python, Java, Go, C, C++.• Deep Learning: CNN, RNN, LSTM, (Variational)AE, Transformer, text embeddings.• ML + Modeling: Tensorflow, PyTorch, LangChain, NLTK.• Cloud + Infrastructure + Scaling: Docker, Kubernetes, AWS, Spark, Airflow.• Interests: AI, Machine Learning, NLP, Information Networks, Graph Theory, AI Safety.
INDUSTRY EXPERIENCE	<div><div>Senior Software Engineer, Amazon Web Services</div><div>Jan. 2022 – now</div><ul style="list-style-type: none">• Optimized system design for Edge devices: Modeled how our current system performs under high-load and failure scenarios (data collection, discrete event simulation, probabilistic modeling, etc). One proposed architecture will increase customer-usable capacity by 20% on identical hardware and while maintaining the same durability targets. This work also served to de-risk several high-potential-impact projects and re-aligned our 2-year product roadmap around efficiency and durability goals.• Predictive health monitoring: Led a team that exposed near-real-time health and performance metrics (including predictive indicators), and designed and built essential automation tooling. This work enabled us to streamline operations for hundreds of AWS Outpost deployments: it directly eliminated 70 hours / week of ops effort, and in the first 3 months, provided early warning about 7 incidents that were likely to escalate to customer-facing issues under the team's previous operational posture.</div> <div><div>Co-founder, Lead ML Engineer, Migacore Technologies</div><div>May 2017 – Dec. 2021</div><ul style="list-style-type: none">• Exit: Acquired by Cirium in March 2021.• Overview: Migacore improves travel demand forecasting by augmenting traditional timeseries forecasts with context gathered around the web, such as information about industry conferences, leisure travel sentiment, and even epidemiological updates. Our partner airlines saw a revenue increase of 1-2% from deploying our signals into their pricing engine.• AI/ML: I designed the end-to-end data acquisition and model training pipeline to be adaptable and interpretable. To briefly summarize: We used NLTK and spaCy to parse natural language webpages into meaningful structured objects. We used off-the-shelf translation tools to ingest non-English text. We trained a semi-supervised model to eliminate irrelevant entries, filtering from about 1B to a more manageable 10M candidates. We trained custom sentiment models with xgboost. We used Pandas (and Dask) to analyze and preprocess ground-truth demand data, combining prices, searches, and bookings data from both individual airlines and third-party booking platforms. We used fb-prophet to make sense of the periodicity in the data. We used Tensorflow to train several overall demand model, including convolutional (CNN) and recurrent (LSTM) variants. We also designed and trained a custom attribution model to help disambiguate which of several concurrent signals was most likely responsible for a single anomaly in observed travel demand (basic idea: use signal similarity as a regularizer; advanced version: use Expectation-Maximization to learn a better similarity function).• Engineering: Our distributed crawling and model training pipeline ran on a dynamically scaled cluster between 10-50 nodes. It scaled to process millions of webpages, and enabled our team to train many models in parallel. We used a microservice architecture to handle different portions of the data collection / preprocessing / labeling / insight generation pipelines.• Team: Recruited, mentored, and grew a team of 8 ML Engineers.</div> <div><div>Software Engineer, Palantir Technologies</div><div>Jun. 2015 – Nov. 2016</div><ul style="list-style-type: none">• Built an internal deployment automation tool, which allows customers to describe their product stack in a custom DSL (<i>Java, Groovy</i>).• Built features on AtlasDB, which is a transactional layer on top of a generic key-value store (<i>Java, Cassandra, distributed systems</i>).</div>

INTERNSHIPS

- **Google** (*New York, 2014*): Built a document classification system to solve a (non-public) internal task that required high interpretability.
- **American Express** (*New York, 2013*): Recommendation system (cardholders & merchants), cold start (non-home city).
- **Groupon** (*San Jose, 2012*): Improved daily deal recommendations, leading to a 3.3% increase in revenue as validated by A/B test.
- **Bloomberg L.P.** (*New York, 2012*): Redesigned the Fair Value Detail function on the Bloomberg terminal.
- **Bloomberg L.P.** (*New York, 2011*): Designed and built an analysis tool for Bloomberg Tradebook's Pair Trading algorithm.

PUBLICATIONS

- Honglei Zhuang, Jing Zhang, **George Brova**, Jie Tang, Hasan Cam, Xifeng Yan, Jiawei Han: "Mining Query-Based Subnetwork Outliers in Heterogeneous Information Networks", ICDM 2014.
- Huan Gui, Yizhou Sun, Jiawei Han, **George Brova**: "Modeling Topic Diffusion in Multi-Relational Bibliographic Information Networks", CIKM 2014.
- Fangbo Tao, **George Brova**, Jiawei Han, Heng Ji, Chi Wang, Brandon Norick, Ahmed El-Kishky, Jialu Liu, Xiang Ren, Yizhou Sun: "NewsNetExplorer: automatic construction and exploration of news information networks", SIGMOD Conference 2014.
- Fangbo Tao, Xiao Yu, Kin Hou Lei, **George Brova**, Xiao Cheng, Jiawei Han, Rucha Kanade, Yizhou Sun, Chi Wang, Lidan Wang, Tim Weninger: "Research-insight: providing insight on research by publication network analysis", SIGMOD Conference 2013.
- Aris Anagnostopoulos, **George Brova**, Evimaria Terzi: "Peer and authority pressure in information-propagation models", Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD) 2011.

EDUCATION

University of Illinois at Urbana-Champaign, Urbana, IL

Sep. 2012 - May. 2015

(Currently on leave)

Ph.D. in Computer Science

Area of study: Applied Machine Learning, Graph Mining

Advisor: Prof. Jiawei Han

Boston University, Boston, MA

Sep. 2008 - May 2012

B.A./M.A. in Computer Science, *Summa Cum Laude*.

GPA: 3.95/4.00 in Computer Science, 3.85/4.00 overall

Trustee Scholar (recipient of a merit-based, four-year, full tuition scholarship)

Dean's List (all semesters)